EXHIBIT 4

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

PALO ALTO NETWORKS, INC.
Petitioner

v.

JUNIPER NETWORKS, INC.
Patent Owner

CASE IPR2013-00369
Patent 7,107,612

DECLARATION OF KEVIN C. ALMEROTH

Dated: March 28, 2014

Respectfully submitted,

Kevin C Almeroth

Kevin Almeroth

DECLARATION OF KEVIN C. ALMEROTH
REGARDING VALIDITY OF THE '612
PATENT

CONFIDENTIAL ATTORNEY EYES ONLY INFORMATION

claimed inventions, and (5) explaining why a person of ordinary skill in the relevant field would have combined the elements in the same manner as in the challenged patent claim.

38. I have also been informed that the claimed invention must be considered as a whole in analyzing obviousness or nonobviousness. In determining the differences between the prior art and the claims, the question under the obviousness inquiry is not whether the differences ***themselves*** would have been obvious, but whether the claimed invention ***as a whole*** would have been obvious. Relatedly, I understand that it may be appropriate to consider whether there is evidence of a "teaching, suggestion, or motivation" to combine the prior art teachings in the prior art, the nature of the problem or the knowledge of a person having ordinary skill in the art.

39. I understand that one indicator of nonobviousness is when prior art "teaches away" from combining certain known elements. For example, a prior art reference teaches away from the patent's particular combination if it leads in a different direction or discourages that combination, recommends steps that would not likely lead to the patent's result, or otherwise indicates that a seemingly inoperative device would be produced.

40. I further understand that certain objective indicia can be important evidence regarding whether a patent is obvious or nonobvious, including commercial success, copying, and industry acceptance or praise. Evidence of such objective indicia must be considered when present. It is generally error to reach a conclusion on obviousness before considering the evidence of secondary considerations, and in then evaluating the latter solely in terms of whether it may fill any gaps in the initial conclusion on obviousness. On the other hand, such evidence is not a requirement for patentability, and the absence of such evidence is a neutral factor in the analysis of obviousness or nonobviousness.

41. I also understand that, in performing a proper unpatentability analysis, an expert must do more than simply provide quotes from the evidentiary record along with conclusory allegations of unpatentability. To the contrary, an expert's conclusions regarding unpatentability must be supported by actual analysis and reasoning set forth in the expert report, such that the theoretical and factual foundation for the expert's conclusions can be properly evaluated.

V. CLAIM CONSTRUCTION

42. It is my understanding that the Board has not finally construed the claim terms for the patents-in-suit in this proceeding. I understand that the Board preliminarily addressed certain claim construction points in its decision to institute the IPR, on which no Juniper expert was permitted testimony. I further understand that the parties have exchanged constructions in a litigation that also addresses the '612 patent.

43. I further understand that the parties agreed that, for purposes of the '612 patent claims, "rules" exist across multiple sessions. *See* Draft Joint Claim Construction Statement transmitted from Juniper to PAN on March 27, 2013. It is my understanding that terms should be given their broadest reasonable construction in an IPR. Under this standard, the terms should

be given their ordinary and customary meaning to one of ordinary skill in the art at the time of invention, unless the patent teaches of a different meaning within the specification.

A. “Rule”

44. As used in the ‘612 patent, a rule must persist across multiple sessions. I understand that all experts who have opined on this term agree that a “rule” in the context of the claims of the ‘612 patent must exist across multiple sessions. I also agree with this construction for the reasons noted below.

45. As noted by the Board, the ‘612 patent describes a rule as a “control policy for filtering incoming and outgoing information packets.” However, this is not the only feature of the claimed rules described by the specification. The ‘612 patent specification, prosecution history, and claims themselves require that rules must exist across multiple sessions. The ‘612 patent sets forth a broad understanding of what constitutes a “rule.” The context of the term as used in the claims makes clear the purpose of “rules” in the invention: they are “for controlling access to and from a network device for incoming and outgoing data packets.” ‘612 patent at 7:48-51 (claim 1). This means that rules contemplate actions to be applied against packets, as in a set of entries for blocking packets from particular source IP addresses. Id. at 5:55-59 (system “allows some packets . . . and denies or drops others” based on rules with “matching criteria” such as “source and destination IP address”); see also 2:61-65; Markman Order (describing rules as involving “actions to be applied against packets”). The ‘612 patent does not require that “rules” be formatted in any particular way or stored in any particular type of data structure.

46. The ‘612 patent does impose one important constraint with respect to this claim term: “rules” are consistently distinguished from other data pertaining solely to a single particular session (i.e., a set of related packets corresponding to a “current application or service”). ‘612 patent at 5:20. Indeed, the ‘612 patent repeatedly identifies an important difference between the use of “rules” as opposed to session-specific data: “[T]he firewall engine may first check a stored look-up table with criteria relating to ongoing current applications or services, before searching the rules.” Id. at 5:14-16; see also id. at 5:51-42 (“current application” data consulted “instead of a rule search”).

47. The ‘612 patent provides an example of how this architecture works, in the context of an FTP session. Session data will be consulted instead of rules if a packet received “is an FTP packet for an FTP [session] that is ongoing.” Id. at 5:21-22. In other words, this approach set forth in the ‘612 patent contemplates that treatment of packets in any single, ongoing FTP session will be handled using session data instead of rules. In this manner, the ‘612 patent architecture makes it possible for “packets in the current application [FTP]” to be handled using the efficient mechanism of a session table lookup “instead of a rule search.” Id. at 5:51-42. The architecture described in the ‘612 patent involves the use of a set of rules in conjunction with a separate data structure referred to as a “flow table” or “session table.” ‘612 patent at 5:14-60. Such a table can keep track of data “corresponding to each current application or service” using

(for example) a common "IP address, port and protocol" for a related group of packets. '612 patent at 5:19-20, 5:37-42. The "current application" could be packets pertaining to a single web session for an e-commerce transaction, a single flow of streaming music or video media, or some other type of network session. '612 patent at 5:19-20, 5:37-42. Once information relating to processing of a session has been written to a flow table for the first packet of that session, the firewall may simply "look up" that information when it receives subsequent packets in the same session. '612 patent at 5:37-42. This allows for faster processing of subsequent packets in the same session, as the flow table may be used "instead of a rule search." '612 patent at 5:37-42. Thus, one defining feature of entries in a flow or session table (and contrasted with rules) is that they exist for only a single session. In other words, while flow tables entries may come and go as new sessions begin and end, the effective lifetimes for rules are not tied to particular sessions, but rather persist across multiple sessions.

48. I note that PAN's expert Dr. Mitchell likewise confirmed the same understanding of the term "rules." In deposition testimony from the Concurrent Litigation, Dr. Mitchell confirmed his understanding that "a rule is something that exists across multiple sessions," Mitchell Deposition at 210:2-211:6, and even pointed to the same portions of the '612 patent specification as supporting the "across multiple sessions" aspect of the claim term "rules."

49. The District Court in the Concurrent Litigation made similar observations regarding the term "rules" in its Markman Order. The Court noted first that the parties had agreed that a "rule" must exist "across multiple sessions." Markman Order at 23. The Court then went on to find that "rules" as contemplated in the '612 patent were distinct from a look-up table data structure, which is used "to describe flow tables" in the '612 patent. See id. at 23 & n.16. In other words, unlike an entry in a flow table or session table, which is deleted following the end of the session, "rules" in the '612 patent are designed to persist across multiple sessions.

50. In light of the foregoing, the broadest reasonable construction of "rules" in this proceeding should include the fundamental concept that rules "exist across multiple sessions." For example, if the Board maintains the other aspects of "rules" mentioned in its Institution Decision, the complete construction should be: "control policy that exists across multiple sessions for filtering incoming and outgoing information packets."

VI. TECHNOLOGY BACKGROUND

51. If asked at trial or at any hearing, I may provide a tutorial regarding technological topics that may be helpful as background. By way of example, these topics could include computer networking principles and standards (such as TCP/IP), fundamentals regarding packet-based communications, development and operation of private and public networks (including the Internet), the development and operation of network security products (such as firewalls and intrusion detection systems), computer programming languages, systems, and methods, principles of electronics such as circuits and integrated circuits, industry practices regarding network threats and security, among other topics.

52. To assist in my testimony, I may rely on sources with which those of ordinary skill in the art would be familiar, including treatises, patents, standards documents (e.g., RFCs), and other publicly available documents, as well as my personal knowledge, background, and personal experience in the field.

53. I have also reviewed certain demonstrative exhibits that were used by the parties during litigation. These documents further informed my opinions expressed in this declaration.

54. Additionally, I have worked with others to help prepare some additional demonstrative exhibits to help explain and illustrate certain concepts in this report. These demonstrative exhibits are attached as Exhibit 2093.

55. As basic background, one of the most widely used computer networks is the Internet. The Internet has been around for several decades. Many trace the origins of the Internet to the Arpanet (the Advanced Research Projects Agency Network), which dates back to the late 1960s. While the origins of the Internet were humble, it has grown into a massive, highly sophisticated network for highly complex and highly varied forms of communication. One of the major leaps in the Internet's evolution did not occur until the early 1990s and the sale of the NSFnet Backbone to MCI, spurring commercialization of the Internet and interest in the World Wide Web (WWW). These changes were significant contributors towards the Internet becoming more widely available and usable.

56. Originally useful mainly for the exchange of text documents through email (using the Simple Mail Transfer Protocol, or SMTP) or file exchange (using a protocol like the File Transfer Protocol, or FTP), the Internet has evolved to support more complex data including multiple media types (e.g., pictures, audio, video), hence the concept of "multimedia." Coupled with new and improved delivery capabilities and increased ways of offering information to users, the ways in which the Internet could be used increased dramatically during the 1990s. These factors led to numerous technical innovations in the way data was made available to users.

57. One of the more important capabilities that existed within the Internet was acting as an information repository whereby servers held information and clients would make requests for that information. The Internet was also evolving such that, instead of servers holding important information, it was other users who held the information. In some cases, instead of information stored in documents, it was the users themselves who were the object of contact, for example, in multimedia conferencing. As described in more detail below, an underlying and long-standing challenge in the Internet was identifying the right address to use in contacting other users or servers.

58. Much Internet communication takes place using a client/server paradigm. That is, content servers hold information desired by users. Through their clients, users make requests for this information, and the server responds by providing the requested information. Such a paradigm is used in, for example, the World Wide Web (WWW). In other applications, like email, servers are responsible for accepting, storing, and forwarding email.

59. Two principles upon which applications and the underlying network infrastructure are based are the use of layered communication to break the task of data delivery into more manageable sub-tasks and the use of protocols to establish rules for how data is communicated.

60. Generally, a protocol is a set of rules that define how a set of functions will be performed. Protocols are important within networks since the two sides of a communication must act in the same, predictable way for data to be successfully delivered. For example, the HyperText Transfer Protocol (HTTP) defines both how requests/responses for objects are to be made and the syntax of request/response messages. The way in which data is exchanged is as important as the format of the data when it is exchanged. Called syntax, protocol specifications typically include the way information in a message is formatted. By clearly describing a protocol's communication rules and syntax, ambiguities and errors can be avoided.

61. Protocols are then combined, based on the layer at which each operates, to perform the functions necessary to deliver data between sources and destinations. In many cases, there is one protocol responsible for the functions of not one, but sometimes multiple layers. Each layer and its corresponding protocol perform a set of functions based on widely, but not universally, agreed upon guidelines. As data is prepared for transmission by an application, it is sent through a set of layers. Each layer performs specified functions. For some of the layers, there is a corresponding protocol and a corresponding protocol header that is added to the application's data.

62. To help understand the process and give direction to the flow of data, the layers are "stacked" one on the other, from the highest layer (the application layer) to the lowest layer (the physical layer). Data, therefore, flows "down" the stack from the application layer of the transmitting host, across the network, and "up" a corresponding stack at the receiver.

63. Over the years, there have been several efforts to "standardize" the layers and the functions performed by each. One example is the International Standards Organization's (ISO) Open System Interconnect (OSI). The OSI stack has seven layers and the general functions of each layer are well-known. ISO's OSI stack model is an older example dating back to the mid-1980s. A more recent example is the "TCP/IP stack," also called the "Internet stack." It integrated the functionality of two of the layers from the OSI stack (Presentation and Session Layers) into the Application Layer and better maps to the Internet's currently used protocols, e.g., IP, UDP, and TCP.

64. Of the layers in the TCP/IP stack, the "highest" layer is the application layer and includes protocols like the HyperText Transfer Protocol (HTTP) and the Simple Mail Transfer Protocol (SMTP). There are dozens of application layer protocols, each typically corresponding to a specific application.

65. The next layer is the transport layer. The two most common protocols are the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP). Where the UDP protocol only provides support for "ports," TCP provides better support for reliable data delivery

through acknowledgements, in-order packet delivery, connections, as well as congestion control, and similar to UDP, port numbers. Data sent through the Internet almost always uses one of these two protocols.

66. The next layer down, and the cornerstone of the Internet, is the Internet layer. The corresponding protocol, the Internet Protocol (IP), provides end-to-end delivery. Using IP address and a variety of support protocols (e.g., routing protocols), routers in the Internet are able to choose the next path towards a destination, thereby robustly moving packets closer to their destination. From a lay perspective, the most common transport protocol, TCP, along with IP, form the core of Internet communications. Hence, the Internet's protocols are commonly called "TCP/IP." Yet, there are two other important layers below IP.

67. The Data Link Layer (DLL) and the Physical Layer are often closely coupled. The reason is that the function of the DLL is to move bits across one physical hop of an end-to-end path. A DLL protocol is typically designed for a specific physical medium, though there are often many different protocols that can be used for a given medium. Physical Layer protocols are responsible for converting digital bits into the analog transmission signal specific to the particular medium being used for communication. It is therefore clear why there is a close relationship between a DLL protocol and the Physical Layer: both work for a specific medium and together move data across a single hop along a path from a source to a destination.

68. Often overlooked in the transmission of data is that DLL protocols–and their headers–only survive across a single hop. Once data is delivered across the hop, the DLL layer header is removed, leaving the IP header exposed, and then based on the next hop to the destination, a new DLL protocol header is added-this one specific to the new medium the packet is to traverse. This process is repeated for each hop along the path from a source to a destination.

69. As mentioned above, while the stack concept is a popular metaphor to help understand how network communication occurs, no reference model is perfect, and each serves as a guideline. Protocols, for example, may perform functions of other layers in violation of a particular reference model, and still be accepted as valid protocols. Even in these cases, the abstraction provided by the general principle of layering and abstraction are sufficient to enable data transmission to take place successfully.

70. In client/server based architectures that use a particular protocol to support a specific application, the protocol is usually implemented in both the client and the server. Thus, for example, there is an HTTP client (i.e., a web browser) and an HTTP server. The client and server communicate over the network using additional protocols focused on the actual delivery of data.

71. There are a number of support protocols that assist in the communication of data between sources and destinations. One such protocol is part of the Domain Name System (see, e.g., RFCs 1033 and 1034, both published in November 1987).

72. In its simplest form, DNS takes a host name (e.g., www.cnn.com) and converts it into an IP address. The IP address is then used to reach the named host.

73. DNS provides a number of additional functions beyond host name-to-address mapping including reverse DNS (host address-to-name mapping), virtual host names (a single host with multiple host names), load balancing (multiple hosts all sharing at least one common name), and email server address mapping.

74. DNS and its mapping functions are useful in that users only need to know the name of the host and are not required to remember more cumbersome IP addresses. Applications typically accept either host names or addresses, and if given a host name, will automatically attempt to resolve it into an address.

VII. PAN HAS FAILED TO SHOW INVALIDITY OF U.S. PATENT NO. 7,107,612

75. In my opinion, PAN has not met its burden of showing that the challenged claims of the ‘612 patent are unpatentable. This opinion is based on my review and analysis of the Mitchell declaration, the facts and evidence upon which he purports to rely (including the prior art and other documents cited in the report), as well as additional facts and my background knowledge and experience in the field. It is further my opinion that the challenged claims are patentable.

A. Overview of the Julkunen Reference

76. Generally, Julkunen discusses a packet filter to allow certain types of connections to pass through a firewall during a single session. It summarizes the work of a student performed using a software firewall programmed on a university Linux computer. The student (Julkunen) provides some specific examples of using his “dynamic packet filter” in connection with single FTP or NFS sessions. In each of the examples, the entries are generated based on a single packet, and once generated last only for the duration of a single session.

B. Overview of the Brenton Reference

77. Brenton is a 672-page book with overview of firewall topics in several different areas. Brenton provides a very high level summary of several concepts related to firewalls, most of which PAN and Dr. Mitchell do not contend relate in any way to the ‘612 patent. Brenton provides, for example, separate descriptions of several different specific types of firewalls.

VIII. OBJECTIVE INDICIA OF NON-OBVIOUSNESS

289. My opinions regarding non-obviousness as set forth above are further supported by additional facts and evidence falling within the category of "objective indicia of non-obviousness." I discuss and analyze these objective indicia in this section (to the extent not already considered above) and conclude that they provide strong additional evidence of non-obviousness in this case.

290. I address a number of the objective indicia of non-obviousness in turn below.

A. Copying

291. [redacted]

292. I understand that both Zuk and Mao spent years working for Juniper and its predecessor Netscreen. (For convenience, I may refer to both Juniper and Netscreen collectively as "Juniper.") I understand that Zuk left Juniper to found PAN, and Mao joined him as a founder of PAN shortly thereafter. [redacted]

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

i. Juniper's Products Embody the Challenged Claims of the '612 Patent

297. Moreover, Juniper's products [redacted] embody the challenged claims of the '612 patent. One of the specific Juniper features [redacted] "IP Action" specifically embodies at least claims 4-7 of the '612 patent.

298. I have considered several Juniper technical documents that demonstrate that Juniper's products [redacted] embody the claims of the '612. [redacted]

299. Juniper's operating system is called JUNOS. Juniper acquired NetScreen, which had its own operating system called ScreenOS. When Juniper acquired NetScreen, it combined many of the features of ScreenOS into JUNOS in a project referred to as [redacted]

[redacted] [redacted] the "IP Action" functionality existing in the ScreenOS that was incorporated in JUNOS products. [redacted]

[redacted] [redacted]

302. IP Action is further described in a ScreenOS reference guide that I have considered. This reference guide makes clear that the IP action feature described above has been implemented in the Juniper products [redacted]. For example, the ScreenOS reference guide

describes how to set up the "IP Action" functionality of the devices at pages 213-214 of the ScreenOS Reference Guide.

303. Moreover, this ScreenOS document describes the "IP sweep" and "port scan" attacks that IP action is meant to protect against. For example, an "IP sweep" is depicted below:

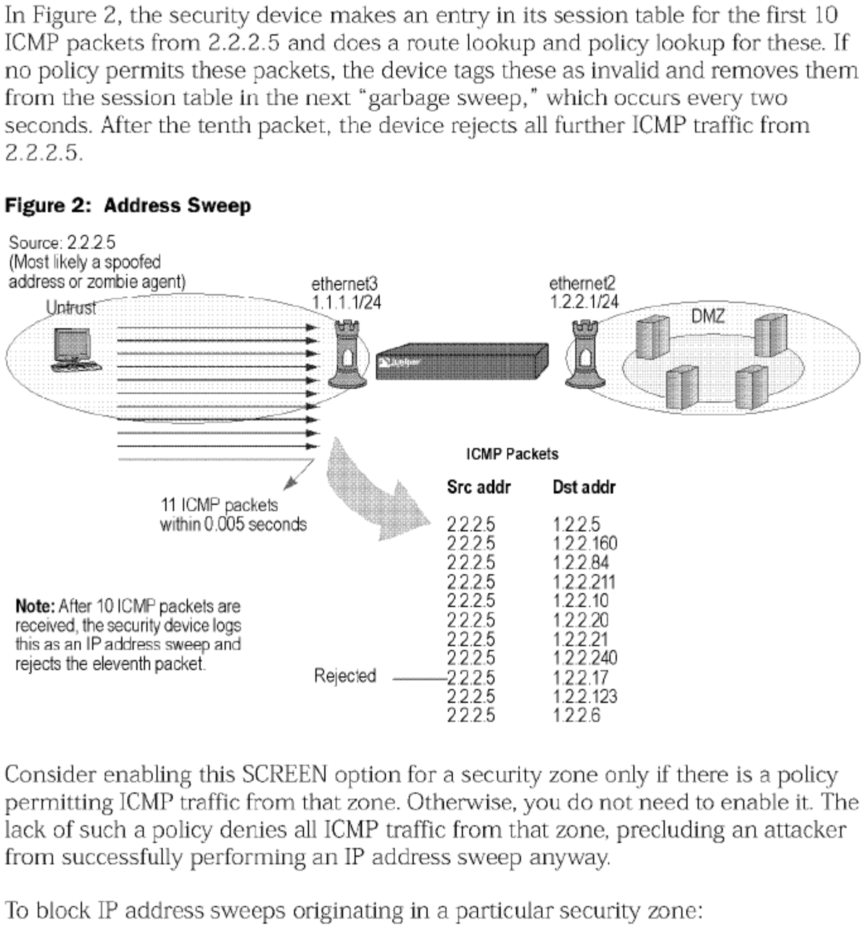

In Figure 2, the security device makes an entry in its session table for the first 10 ICMP packets from 2.2.2.5 and does a route lookup and policy lookup for these. If no policy permits these packets, the device tags these as invalid and removes them from the session table in the next "garbage sweep," which occurs every two seconds. After the tenth packet, the device rejects all further ICMP traffic from 2.2.2.5.

Figure 2: Address Sweep

Note: After 10 ICMP packets are received, the security device logs this as an IP address sweep and rejects the eleventh packet.

ICMP Packets

Src addr	Dst addr
2.2.2.5	1.2.2.5
2.2.2.5	1.2.2.160
2.2.2.5	1.2.2.84
2.2.2.5	1.2.2.211
2.2.2.5	1.2.2.10
2.2.2.5	1.2.2.20
2.2.2.5	1.2.2.21
2.2.2.5	1.2.2.240
2.2.2.5	1.2.2.17
2.2.2.5	1.2.2.123
2.2.2.5	1.2.2.6

Consider enabling this SCREEN option for a security zone only if there is a policy permitting ICMP traffic from that zone. Otherwise, you do not need to enable it. The lack of such a policy denies all ICMP traffic from that zone, precluding an attacker from successfully performing an IP address sweep anyway.

To block IP address sweeps originating in a particular security zone:

304. Port Scan is depicted below:

> In Figure 3, the security device makes an entry in its session table for the first 10 connection attempts from 2.2.2.5 to 1.2.2.5 and does a route lookup and policy lookup for these. If no policy permits these connection attempts, the device tags these as invalid and removes them from the session table in the next "garbage sweep," which occurs every two seconds. After the tenth attempt, the device rejects all further connection attempts from 2.2.2.5.

Figure 3: Port Scan

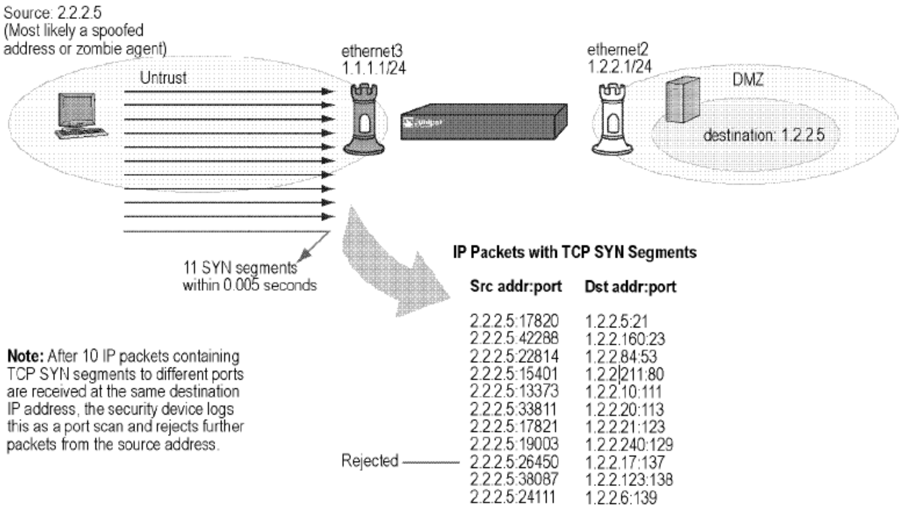
Source: 2.2.2.5
(Most likely a spoofed address or zombie agent)
Untrust
ethernet3
1.1.1.1/24
ethernet2
1.2.2.1/24
DMZ
destination: 1.2.2.5
11 SYN segments within 0.005 seconds
IP Packets with TCP SYN Segments
Rejected

Src addr:port	Dst addr:port
2.2.2.5:17820	1.2.2.5:21
2.2.2.5:42288	1.2.2.160:23
2.2.2.5:22814	1.2.2.84:53
2.2.2.5:15401	1.2.2|211:80
2.2.2.5:13373	1.2.2.10:111
2.2.2.5:33811	1.2.2.20:113
2.2.2.5:17821	1.2.2.21:123
2.2.2.5:19003	1.2.2.240:129
2.2.2.5:26450	1.2.2.17:137
2.2.2.5:38087	1.2.2.123:138
2.2.2.5:24111	1.2.2.6:139

Note: After 10 IP packets containing TCP SYN segments to different ports are received at the same destination IP address, the security device logs this as a port scan and rejects further packets from the source address.

305. Each of these features were incorporated into the later JUNOS operating system, [redacted]. For example, the same figures describing the protection against IP sweeps and Port-scan attacks appear in the Juniper "Security Configuration Guide" for the JUNOS operating system, release 9,2., at pages 157-159.

306. The features of IP Action, [redacted], remain available in JUNOS. For example the JUNOS Security Configuration Guide describes how these features can be configured in the Juniper products:

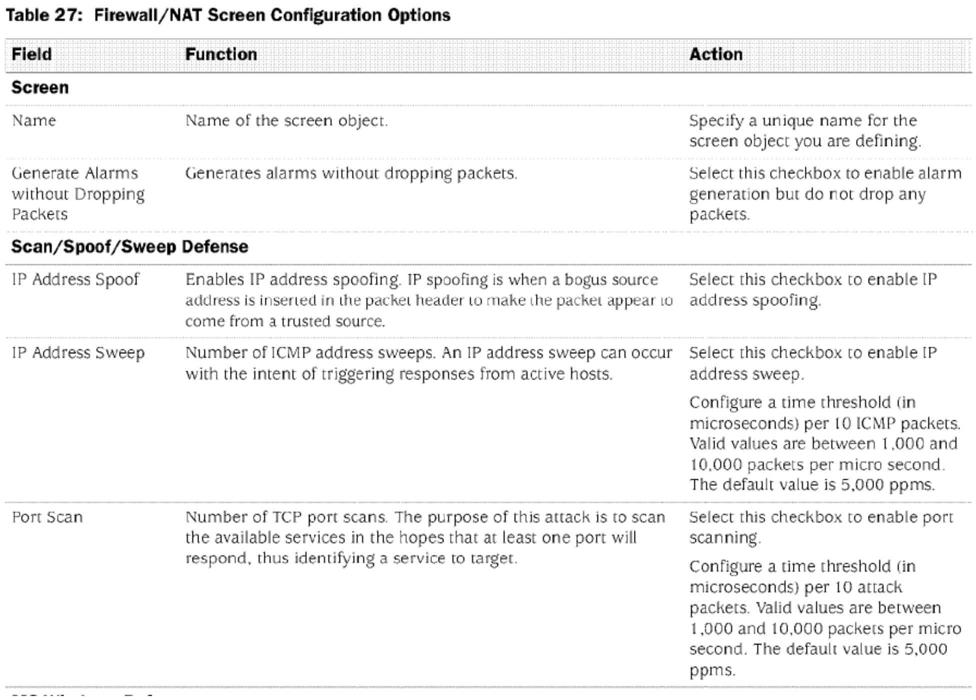

Table 27: Firewall/NAT Screen Configuration Options

Field	Function	Action
Screen		
Name	Name of the screen object.	Specify a unique name for the screen object you are defining.
Generate Alarms without Dropping Packets	Generates alarms without dropping packets.	Select this checkbox to enable alarm generation but do not drop any packets.
Scan/Spoof/Sweep Defense		
IP Address Spoof	Enables IP address spoofing. IP spoofing is when a bogus source address is inserted in the packet header to make the packet appear to come from a trusted source.	Select this checkbox to enable IP address spoofing.
IP Address Sweep	Number of ICMP address sweeps. An IP address sweep can occur with the intent of triggering responses from active hosts.	Select this checkbox to enable IP address sweep. Configure a time threshold (in microseconds) per 10 ICMP packets. Valid values are between 1,000 and 10,000 packets per micro second. The default value is 5,000 ppms.
Port Scan	Number of TCP port scans. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.	Select this checkbox to enable port scanning. Configure a time threshold (in microseconds) per 10 attack packets. Valid values are between 1,000 and 10,000 packets per micro second. The default value is 5,000 ppms.

307. Juniper's products embody the IP Action feature [redacted]. And the IP Action feature, when implemented on a firewall that has NAT capability as many Juniper products do, practices Claims 4-6 of the '612 patent.

B. Other objective indicia

308. Moreover, there are other objective indicia that weigh in favor of a finding of non-obviousness. For example, the Juniper embodying products—which prominently feature the embodying functionality in documentation (see, e.g., Ex. 2072 at 27; Ex. 2071 at 213)—have been very commercially successful. Ex. 2053. These commercially successful products embody claim claims. See § VIII.A.i. [redacted]

[redacted]

network addresses and port numbers associated with the firewall and with corresponding destination nodes in the private network, are supported, for example, by the application's disclosure of a "public network link," "private network link," and network address translation. E.g., id. at 19.

328. I explained above in section V.A that the '612 patent specification discloses rules spanning multiple sessions, as distinguished from session-specific data. Page 16-18 and 30 of the specification as originally filed (Ex. 1005) include the exact material that Dr. Almeroth cites from the '612 issued specification, so the application as originally filed also disclosed rules that exists across multiple sessions.

329. At a high level, the '612 patent describes network security technology that can be used in, for example, a firewall. A firewall can help protect networks in a number of ways by using and applying "rules" to network traffic.

330. The '612 patent acknowledges that using only "a fixed set of rules can be restrictive in many practical applications." '612 patent at 3:1-2. Rules in early firewalls were static, and stayed in place until removed or modified by a network administrator. Thus, there was a "need in the art for a firewall engine which can generate rules ***dynamically***, based upon information extracted from incoming packets . . . ." *Id.* at 3:9-12.

331. To address shortcomings of the prior art, the '612 patent describes "[a]n improved firewall for providing network security" using novel features such as the ability to add or modify rules in a set of rules based on a sequence of data packets received. *Id.* at Abstract. The improved firewall may also "provide[] for translation of IP addresses between the firewall and the internal network." *Id.*

332. As described in the following sections, the substitute claims including this added limitation are not anticipated by, or obvious in view of the closest known prior art.

B. Patentable Distinction From the Closest Known References

333. I have been asked to identify the closet known set of references[3] with respect to the new claims as proposed above. To identify these references, I have considered the references submitted to the PTAB as part of this petition (*i.e.*, Julkunen, Schneider, Brenton, and IETF NAT). These references are addressed in detail earlier in this report. In conducting my review, I noticed at times there were references that used words such as "dynamic" or "rules" in a variety of different ways (including ways that were inconsistent with the invention of the '612). In these cases, what I endeavored to do was to closely study the actual technical details set forth in these

[3] The inclusion of a particular reference in the set of closest related references should not be interpreted as an admission that the reference is prior art. I have not undertaken an analysis to determine whether any of these references are in fact prior art.

references rather than focusing on an author's particular word choice. In so doing, I was able to identify disclosures that used this type of language but that was actually using technology plainly distinct from what was claimed in the '612 patent. For all of the reasons given throughout this report, the proposed amended claims would not be invalidated by any of these references

334. As I understand, the '612 patent has been the subject of litigation. I understand Dr. Mitchell has submitted an expert report on invalidity of the '612 patent. I have further considered all of the references discussed by Dr. Mitchell and included those in the set of closest known references and addressed them below. I have also considered all of the references that PAN has identified in its invalidity contentions in the Delaware Litigation, which is a larger set than what was ultimately considered by Dr. Mitchell in his litigation expert report. I have also considered both PAN's invalidity contentions and Juniper's validity contentions from the Delaware litigation. I have further considered all of the references considered during the prosecution of the '612 patent. I have also undertaken an independent attempt to identify related references to the '612 patent beyond what has been identified above, based on my knowledge and experience in the field.

335. The resulting set of closest known references that I have considered as part of my analysis include[4]:

- Julkunen
- Brenton
- U.S. Patent No. 6,178,505 to Schneider et al. ("Schneider '505 Patent")
- U.S. Patent No. 6,098,172 to Coss et al. ("Coss '172 Patent")
- U.S. Patent No. 6,141,749 to Coss et al. ("Coss '749 Patent")
- U.S. Patent No. 6,496,935 to Fink and Harush ("Fink '935 Patent")
- U.S. Patent No. 6,701,432 to Deng el al. ("Deng '432 Patent")
- U.S. Patent No. 5,968,176 to Nessett and Sherer ("Nessett '176 Patent")
- U.S. Patent No. 6,651,099 to Dietz et al. ("Dietz '099 Patent")

[4] The full set of closest known references I considered in my analysis and submitted herewith are Exs. 2043 – 2072.

- U.S. Patent No. 5,606,668 to Shwed ("Shwed `668 Patent")
- U.S. Patent No. 5,835,726 to Shwed et al. ("Shwed '726 Patent")
- European Patent Application 658,837 to Shwed ("Shwed `837 European Patent Application")
- U.S. Patent No. 5,951,651 to Lakshman et al. ("Lakshman `651 Patent")
- U.S. Patent No. 5,983,270 to Abraham et al. ("Abraham `270 Patent")
- U.S. Patent No. 6,009,475 to Shrader ("Shrader `475 Patent")
- U.S. Patent No. 6,016,310 to Muller et al. ("Muller `310 Patent")
- U.S. Patent No. 6,400,707 to Baum et al. ("Baum `707 Patent")
- U.S. Patent No. 6,757,680 to Choy ("Choy `680 Patent")
- U.S. Patent No. 6,845,452 to Roddy et al. ("Roddy `452 Patent")
- U.S. Patent No. 7,013,482 to Krumel ("Krumel `482 Patent")
- European Patent Application 893,921 to Sheldrick ("Sheldrick `921 European Patent Application")
- U.S. Patent Application No. 2002/0027907 to Tateoka ("Tateoka `907 Patent Application")
- U.S. Patent Application No. 2002/0188720 to Terrell et al. ("Terrell `720 Patent Application")
- D. Decasper, Z. Dittia, G. Parulkar, B. Plattner, "Router Plugins: A Software Architecture for Next Generation Routers" ("Decasper Reference")
- H. Baraka, H. El-Manawy, and A. Attiya, "An Integrated Model for Intranet Security Using Prevention and Detection Techniques" ("Baraka Reference")
- Check Point documentation including "Check Point FireWall-1 Quick Start Guide" (Version 4.0); "Check Point Firewall-1 Architecture and Administration" (Volume 4.0); and/or "Managing Check Point FireWall-1 Using the OpenLook GUI" (Version 4.0) ("Check Point FireWall-1")

- Related references used to make obviousness combinations including IETF RFC 959 entitled “FILE TRANSFER PROTOCOL (FTP)”; IETF RFC 1631 entitled “The IP Network Address Translator”; and IETF Internet Draft draft-ietf-nat-protocol-complications-00.txt entitled “Protocol Complications with the IP Network Address Translator (NAT)

336. None of the closest known references discloses or renders obvious every limitation of the proposed amended claims. In my analysis below, I have focused particularly on identifying limitations from the proposed amended independent claims. As I conclude that none of the closest known references disclose all of the limitations of the proposed amended independent claims, the related dependent claims are also valid for at least the reason that the claim on which each depends is valid.

337. The Abstract of the Schneider ‘505 patent provides as follows:

> A scalable access filter that is used together with others like it in a virtual private network to control access by users at clients in the network to information resources provided by servers in the network. Each access filter uses a local copy of an access control data base to determine whether an access request made by a user. Changes made by administrators in the local copies are propagated to all of the other local copies. Each user belongs to one or more user groups and each information resource belongs to one or more information sets. Access is permitted or denied according to of access policies which define access in terms of the user groups and information sets. The rights of administrators are similarly determined by administrative policies. Access is further permitted only if the trust levels of a mode of identification of the user and of the path in the network by which the access is made are sufficient for the sensitivity level of the information resource. If necessary, the access filter automatically encrypts the request with an encryption method whose trust level is sufficient. The first access filter in the path performs the access check and encrypts and authenticates the request; the other access filters in the path do not repeat the access check.

338. The Schneider ‘505 patent was included as part of the Petition for *Inter Partes* Review, but the Board did not authorize a review using the Schneider ‘505 patent on the grounds that it was redundant with the grounds on which the review was granted. Similar to the Julkunen reference, the Schneider ‘505 patent fails to disclose limitations of either the claims as currently written or the proposed amended claims.

339. The Schneider ‘505 patent does not mention firewalls specifically instead describing an “access filter.” This access filter is more directed at controlling access to information sets by individual users and groups of users. Based on trust relationships and policies established by, for example, an administrator, access attempts are either filtered or

allowed. As such, the Schneider '505 patent does not describe the creation of dynamic rules that last beyond the confines of a session and are based on the inspection of a sequence of packets. For at least these reasons, the Schneider '505 patent does not disclose each and every limitation of any of the proposed amended claims of the '612 patent.

340. The Abstract of the Coss '172 patent provides as follows:

> Computer network firewalls which include one or more features for increased processing efficiency are provided. A firewall in accordance with the invention can support multiple security policies, multiple users or both, by applying any one of several distinct sets of access rules. The firewall can also be configured to utilize "stateful" packet filtering which involves caching rule processing results for one or more packets, and then utilizing the cached results to bypass rule processing for subsequent similar packets. To facilitate passage to a user, by a firewall, of a separate later transmission which is properly in response to an original transmission, a dependency mask can be set based on session data items such as source host address, destination host address, and type of service. The mask can be used to query a cache of active sessions being processed by the firewall, such that a rule can be selected based on the number of sessions that satisfy the query. Dynamic rules may be used in addition to pre-loaded access rules in order to simplify rule processing. To unburden the firewall of application proxies, the firewall can be enabled to redirect a network session to a separate server for processing.

341. As can be seen from the Abstract, the focus of the Coss '172 patent is on improving the efficiency of a firewall so it can operate on faster links without introducing additional delay. The focus is on caching rule-processing results for "similar packets" in a particular "active session," and therefore, is ultimately directed to different subject matter than the claims of the '612 patent.

342. The Coss '172 patent was involved in prior litigation, and in that litigation, Dr. Mitchell offered an opinion that the alleged dynamic rules he identified lasted only for the life of a session. Ex. 2094 ¶ 729. To the extent anything in the Coss '172 patent includes any firewall constraint that may last longer than a session, they relate to other "aspects" of the firewall and not to the "fourth aspect," *i.e.,* the part of the firewall that uses a received packet to create a flow table entry. (*see, e.g.,* 2:30-31).

343. The Coss '172 patent also does not disclose receiving a sequence of packets. The one citation relied upon by Dr. Mitchell in the previous litigation (*i.e.,* 7:8-9) describes reception of only a single IP packet. *i.e.,* Ex. 2063 at 7:8-9

344. The Coss '749 patent includes a similar, if not identical disclosure, which likewise fails to anticipate the '612 patent for the same reasons set forth above.

345. For at least the reasons above, the Coss '172 patent and Coss '749 patent to not disclose each and every limitation of any of the proposed amended claims of the '612 patent.

346. The Abstract of the Fink '935 patent provides as follows:

> A system, a device and a method for accelerating packet filtration by supplementing a firewall with a pre-filtering module. The pre-filtering module performs a limited set of actions with regard to the packets, according to whether the packets are received from a connection which has been previously permitted by the firewall. If the packets are received from such a permitted connection, then the pre-filtering module forwards the packets to their destination, optionally performing one or more actions on the packets. Otherwise, the packets are forwarded to the firewall for handling. Preferably, once the firewall has transferred responsibility for the connection to the pre-filtering module, or "off-loaded" the connection, the firewall does not receive further packets from this connection until a timeout occurs for the connection, or a packet is received with particular session-control field values, such that the connection is closed. Optionally and preferably, the pre-filtering module is implemented as hardware.

347. The focus of the Fink '935 patent is on developing a pre-filtering module with the goal of improving efficiency and reducing the processing burden of handling many packets per second. And while the Fink '935 patent does use what it describes as rules to determine whether packets should be allowed through the firewall, the Fink '935 patent does not describe the creation of dynamic rules that last beyond the confines of a session and are based on the inspection of a sequence of packets. For at least these reasons, the Fink '935 patent does not disclose each and every limitation of any of the proposed amended claims of the '612 patent.

348. The Abstract of the Deng '432 patent provides as follows:

> A gateway for screening packets transferred over a network. The gateway includes a plurality of network interfaces, a memory and a memory controller. Each network interface receives and forwards messages from a network through the gateway. The memory temporarily stores packets received from a network. The memory controller couples each of the network interfaces and is configured to coordinate the transfer of received packets to and from the memory using a memory bus. The gateway includes a firewall engine coupled to the memory bus. The firewall engine is operable to retrieve packets from the memory and screen each packet prior to forwarding a given packet through the gateway and out an appropriate network interface. A local bus is coupled between the firewall engine and the memory providing a second path for retrieving packets from memory when the memory bus is busy. An expandable external rule memory is coupled to the local bus and includes one or more rule sets accessible by the firewall engine using the local bus. The firewall engine is operable to retrieve rules from a rule set and screen packets in accordance with the retrieved rules.

349. The focus of the Deng '432 patent is focused on an architecture that provides multiple bus interfaces and better memory management in order to improve the efficiency of packet processing and reduce delays associated with handling large numbers of packets per second. And while the Deng '432 patent does use what it describes as rules to determine whether packets should be allowed through the firewall, the Deng '432 patent does not describe the creation of dynamic rules that last beyond the confines of a session and are based on the inspection of a sequence of packets. For at least these reasons, the Deng '432 patent does not disclose each and every limitation of any of the proposed amended claims of the '612 patent.

350. The Abstract of the Nessett '176 patent provides as follows:

> A system provides for establishing security in a network that include nodes having security functions operating in multiple protocol layers. Multiple network devices, such as remote access equipment, routers, switches, repeaters and network cards having security functions are configured to contribute to implementation of distributed firewall functions in the network. By distributing firewall functionality throughout many layers of the network in a variety of network devices, a pervasive firewall is implemented. The pervasive, multilayer firewall includes a policy definition component that accepts policy data that defines how the firewall should behave. The policy definition component can be a centralized component, or a component that is distributed over the network. The multilayer firewall also includes a collection of network devices that are used to enforce the defined policy. The security functions operating in this collection of network devices across multiple protocol layers are coordinated by the policy definition component so that particular devices enforce that part of the policy pertinent to their part of the network.

351. The focus of the Nessett '176 patent is on developing a coordinated security policy across a complex multi-layer organization. One goal of the Nessett '176 patent is to develop a policy definition framework such that efforts across the organization are coordinated and holes in the security apparatus are not created. And while the Nessett '176 patent does use what it describes as rules to determine whether packets should be allowed through various firewalls, the Nessett '176 patent does not describe the creation of dynamic rules that last beyond the confines of a session and are based on the inspection of a sequence of packets. For at least these reasons, the Nessett '176 patent does not disclose each and every limitation of any of the proposed amended claims of the '612 patent.

352. The Abstract of the Dietz '099 patent provides as follows:

> A monitor for and a method of examining packets passing through a connection point on a computer network. Each packets conforms to one or more protocols. The method includes receiving a packet from a packet acquisition device and performing one or more parsing/extraction operations on the packet to create a parser record comprising a function of selected portions of the packet. The parsing/extraction operations depend on one or more of the protocols to which the

> packet conforms. The method further includes looking up a flow-entry database containing flow-entries for previously encountered conversational flows. The lookup uses the selected packet portions and determining if the packet is of an existing flow. If the packet is of an existing flow, the method classifies the packet as belonging to the found existing flow, and if the packet is of a new flow, the method stores a new flow-entry for the new flow in the flow-entry database, including identifying information for future packets to be identified with the new flow-entry. For the packet of an existing flow, the method updates the flow-entry of the existing flow. Such updating may include storing one or more statistical measures. Any stage of a flow, state is maintained, and the method performs any state processing for an identified state to further the process of identifying the flow. The method thus examines each and every packet passing through the connection point in real time until the application program associated with the conversational flow is determined.

353. The focus of the Dietz '099 patent is on network monitoring and deep packet inspection. The goal of the reference is to describe the mechanism by which each packet is inspected and determined to be associated with a particular flow. While the Dietz '099 patent does not mention firewalls specifically or prohibiting packets from entering a network, the Dietz '099 patent does describe concepts related to packet inspection. Further, the Dietz '099 patent does not describe the creation of dynamic rules that last beyond the confines of a session and are based on the inspection of a sequence of packets. For at least these reasons, the Dietz '099 patent does not disclose each and every limitation of any of the proposed amended claims of the '612 patent.

354. The Abstract of the Decasper Reference provides as follows:

> Present day routers typically employ monolithic operating systems which are not easily upgradable and extensible. With the rapid rate of protocol development it is becoming increasingly important to dynamically upgrade router software in an incremental fashion. We have designed and implemented a high performance, modular, extended integrated services router software architecture in the NetBSD operating system kernel. This architecture allows code modules, called plugins, to be dynamically added and configured at run time. One of the novel features of our design is the ability to bind different plugins to individual flows; this allows for distinct plugin implementations to seamlessly coexist in the same runtime environment. High performance is achieved through a carefully designed modular architecture; an innovative packet classification algorithm that is both powerful and highly efficient; and by caching that exploits the flow-like characteristics of Internet traffic. Compared to a monolithic best-effort kernel, our implementation requires an average increase in packet processing overhead of only 8%, or 500 cycles/2.1ms per packet when running on a P6/233.

355. The focus of the Decasper Reference is on building an open-source, modular router such that existing components can be refined and enhanced by other

researchers/developers and new modules can be incorporated into the routing architecture. While one of the modules described is a firewall, there is little detail as to what functions the firewall will perform and how they are to be performed. One of skill in the art reading the paper would understand that the firewall functionality disclosed would have been that typical of firewalls at the time the paper was published. While there is a description of the firewall module analyzing packets, grouping them into flows, and applying different policies to different flows, the Decasper Reference does not describe the creation of dynamic rules that last beyond the confines of a session and are based on the inspection of a sequence of packets. For at least these reasons, the Decasper Reference patent does not disclose each and every limitation of any of the proposed amended claims of the ‘612 patent.

356. The Abstract of the Baraka Reference provides as follows:

> This paper presents a new model for securing an Intranet, connected to the Internet, based on a hybrid technique. The model integrates two security modules; the Prevention module and the Detection module. The proposed model provides a dynamic binding between the two modules. A comparison between the proposed model and the classical security techniques proved the effectiveness of the new model.

357. The general focus of the Baraka reference is to describe a model for combining intrusion detection technology with a firewall. It provides a number of examples of possible applications for this technology but without describing specifics of how it is to be implemented. Indeed, the text of Baraka is less than five pages and lacks significant technical detail. This disclosure falls short of describing a set of rules to which dynamic rules may be added or modified, which last beyond the confines of a session, and are based on the inspection of a sequence of packets. For at least these reasons, the Baraka reference does not disclose each and every limitation of any of the proposed amended claims of the ‘612 patent. Moreover, Baraka has no publication date, and no indication that is was published at all, let alone before the critical date.

358.

359. There are a set of three references that relate to the Check Point Firewall-1 product including “Check Point FireWall-1 Quick Start Guide” (Version 4.0); “Check Point Firewall-1 Architecture and Administration” (Volume 4.0); and “Managing Check Point FireWall-1 Using the OpenLook GUI” (Version 4.0). Together or individually, these references describe a traditional firewall, including a management module:

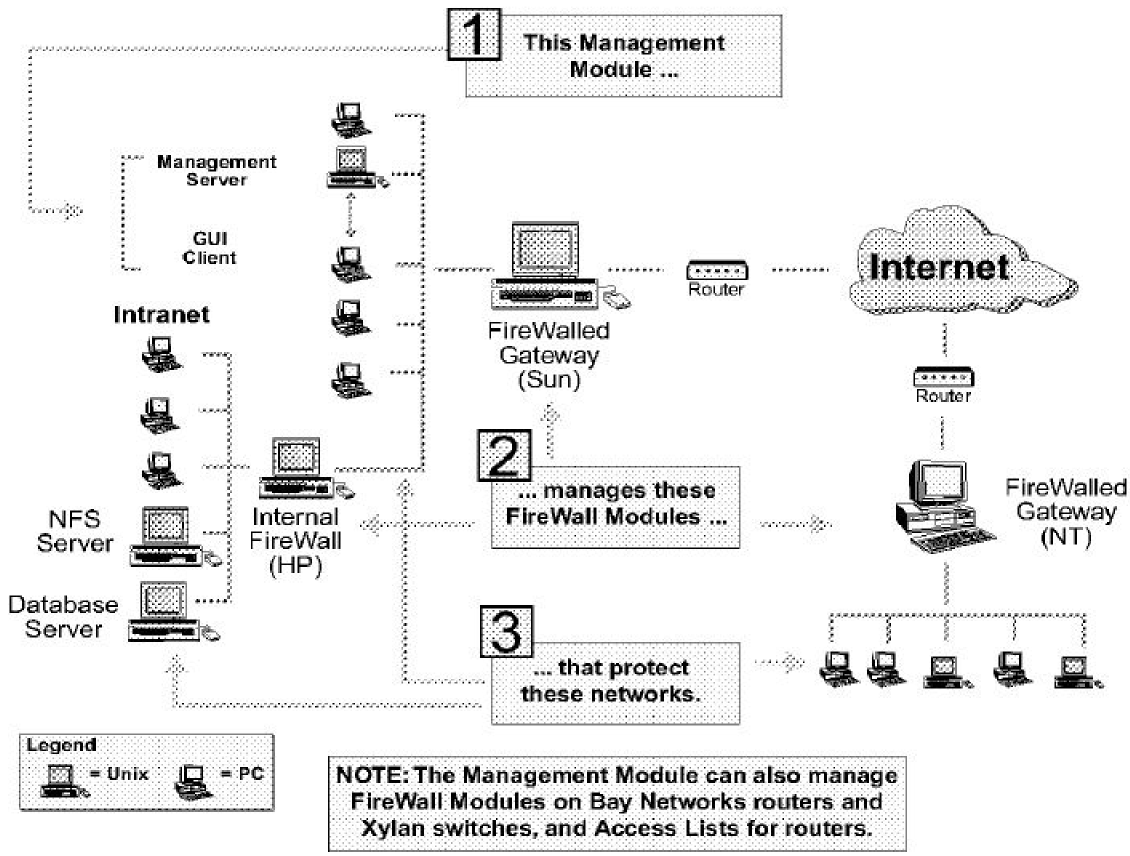
1
This Management Module ...
Management Server
GUI Client
Intranet
FireWalled Gateway (Sun)
Router
Internet
Router
2
... manages these FireWall Modules ...
FireWalled Gateway (NT)
NFS Server
Internal FireWall (HP)
Database Server
3
... that protect these networks.
Legend
= Unix
= PC
NOTE: The Management Module can also manage FireWall Modules on Bay Networks routers and Xylan switches, and Access Lists for routers.

FIGURE 1-1 Distributed FireWall-1 Configuration

See PAN000526282.

360. The Check Point Firewall-1 product includes what it describes as "rules," but these are not the rules of the claims, i.e., they are not dynamically generated in response to data extracted from a sequence of data units. And while there is a separate mechanism for dealing with TCP SYN attacks, this mechanism (*i.e.*, SYNDefender Gateway or SYNDefender Passive Gateway) does not result in the creation of rules used by the firewall. Therefore, Check Point Fire Wall-1 does not describe the creation of dynamic rules that last beyond the confines of a session and are based on the inspection of a sequence of packets. For at least these reasons, Check Point Fire Wall-1 does not disclose each and every limitation of any of the proposed amended claims of the '612 patent.

361. The Abstract of the Shwed '668 patent provides as follows:

> A filter module allows controlling network security by specifying security rules for traffic in the network and accepting or dropping communication packets according to these security rules. A set of security rules are defined in a high level form and are translated into a packet filter code. The packet filter code is loaded

> An apparatus comprising a network interface, through which the apparatus facilitates communication between a client device and a remote device and a controller is presented. In accordance with one aspect of the present invention, the controller, coupled to the network interface, dynamically creates and removes admission filters based, at least in part, on an admissions profile that, when triggered, the filter(s) initiate an admission control decision preventing premature allocation of resources which are not used or authorized.

385. The focus of the Terrell '720 Patent Application is a straightforward reference that describes the dynamic creation of filters last the duration of a session only and are developed based on a decision whether the flow should be allowed into the protected system. Therefore, the Terrell '720 Patent Application does not describe the creation of dynamic rules that last beyond the confines of a session and are based on the inspection of a sequence of packets. For at least these reasons, the Terrell `720 Patent Application does not disclose each and every limitation of any of the proposed amended claims of the `612 patent.

386. There are several additional references that have used as obviousness references in combination with other primary references. These include, at least the FTP RFC, the NAT RFC, and the NAT Complications IETF RFC 959 entitled "FILE TRANSFER PROTOCOL (FTP)"; IETF RFC 1631 entitled "The IP Network Address Translator"; and IETF Internet Draft draft-ietf-nat-protocol-complications-00.txt entitled "Protocol Complications with the IP Network Address Translator (NAT). These references have been proposed in various combinations in an attempt to meet certain limitations of dependent claims. None of these references discuss firewalls generally or specifically and have not been suggested for any of the limitations of the independent claims. For the sake of completeness, therefore, it is clear that, individually, they do not meet any or all of the limitations of the proposed amended claims of the '612 patent.

387. Prior to the invention of the '612 patent, firewalls commonly used a fixed set of rules in performing network security functions. Ex. 1001, at 3:1-2. Rules in early firewalls were static, and stayed in place until removed or modified by a network administrator. Later firewalls introduced "flow-based" functionality that allowed for greater efficiency by making firewall decisions for the first packet in a "session" of related packets, and then applying that decision to other packets in the same session or flow. Ex. 1001, at 5:50-55. These systems still did not permit dynamic flexibility in responding to network traffic in a manner that would extend beyond a single session; the static rules in the system persisted across multiple sessions and as with older firewalls these rules would not be altered absent human intervention.

388. None of the closest known prior art anticipates or renders obvious the proposed substitute claims.

389. Moreover, the objective indicia discussed in detail above in Section VIII, above, also applies to the proposed substitute claims for all of the same reasons discussed.

390. There is no reference that suggests the combination of, for example: 1) dynamic creation 2) of rules that persist across multiple sessions 3) based on the inspection of a sequence of data units.

391. Neither Julkunen, Brenton, nor any of the closest known prior art either discloses or renders obvious the concept of adding or modifying rules that exist across multiple sessions based on data extracted from multiple incoming packets.

392. I considered the knowledge of those having ordinary skill in the art, specifically with respect to the features of the amended claims that provide the basis of patentable distinction. Such individuals had an understanding of the state of the art of network security, including knowledge of fundamental firewall technology that developed in response to the security concerns that arose in response to increasing numbers of individuals were connecting to the Internet. It was well-understood that network security systems were generally managed by network administrators capable of configuring sets of rules in a firewall to guard against anticipated attacks. These sets of rules would sometimes become quite complex, and required detailed knowledge of the configuration of the network. After carefully configuring these rules, they would be implemented (or "committed") in the firewall and fixed in place until the network administrator needed to make further adjustments.

393. Over time, those of skill in the art incorporated optimizations to these firewalls. For example, engineers in the field realized that all packets in a single session can often be treated in the same way. Thus, many firewalls became "flow-based" (or "session-based") and incorporated session-specific data for packets in a flow table or session table. *see also* Ex. 2065. Notwithstanding that these flow-based firewalls had some ability to actively respond to new sessions, the firewall rules remained fixed across multiple sessions.

394. Thus, at the time of the '612 patent invention (and as demonstrated in the closest known prior art), the virtually ubiquitous network security paradigm was one in which firewalls had fixed rules that persisted over time, alongside a session table to which session-specific entries could be added and deleted as sessions were initiated and terminated. Those of skill in the art saw this architecture as a way of balancing competing needs of flexibility, speed, and security. Network security engineers also tried to improve upon these various trade-offs in existing systems. For example, as shown above, the prior art disclosed methods for facilitating rule updates by network administrators, introduced approaches for efficient handling of session-based data, and attempted to optimize aspects of the hardware and software in existing firewall implementations to increase processing speed.

395. One of the insights of the '612 patent—not disclosed in or rendered obvious by any of the preexisting art—was that ***dynamic capabilities could be built into the rule set itself,*** as opposed to a session table or similar functionality limited to a single session. This contravened the standard approach where it was seen as important to have rules fixed and unchanging.

Indeed, those skilled in the art generally believed that, due to complexity of existing rule bases, "committing" rules was something that a human administrator had to oversee.

396. As set forth in greater detail in the Patent Owner Response, the '612 patent's approach of adding rules dynamically without user intervention proved to be very successful. The feature was incorporated into many successful NetScreen and Juniper products. Moreover, competitors such as PAN [redacted] These facts further demonstrate the novelty and non-obviousness of the '612 patent claims as amended.